

A Counter Example to the Theory of Simultaneous Localization and Map Building

Simon J. Julier

Virtual Reality Laboratory
Code 5580

ITT Systems/Naval Research Laboratory
Washington DC 20375-5320
julier@ait.nrl.navy.mil

Jeffrey K. Uhlmann

Advanced Information Technology
Code 5580

Naval Research Laboratory
Washington DC 20375-5320
uhlmann@ait.nrl.navy.mil

Abstract

This paper analyses the properties of the full covariance simultaneous map building problem (SLAM). We prove that, for the special case of a stationary vehicle (with no process noise) which uses a range-bearing sensor and has non-zero angular uncertainty, the full-covariance SLAM algorithm always yields an inconsistent map. We also show, through simulations, that these conclusions appear to extend to a moving vehicle with process noise. However, these inconsistencies only become apparent after several hundred beacon updates.

1 Introduction

Autonomous Guided Vehicle (AGV) technology plays a prominent role in a wide variety of scientific, industrial, and military applications. In all of these applications there is a critical need for *localization* — the AGV must be able to accurately and repeatedly estimate its own position. The most general localization strategy is to equip the AGV with the capability to construct its own map of beacons (which can be artificial or naturally-occurring). The AGV maintains a continuous estimate of its absolute position as it re-observes the dynamically mapped beacons. This strategy is referred to as Simultaneous Localization And Map building (SLAM).

The seminal work on the rigorous application of the Kalman filter to the SLAM problem was carried out in the late 1980s by Smith, Self and Cheeseman [1] who introduced the notion of a *stochastic map*. Previous researchers addressed the SLAM problem by assuming that the beacon and vehicle estimates could be propagated independently of one another. What [1] showed was that the vehicle and beacon estimates are not independent of one another and a complete joint covariance matrix comprising the vehicle and all beacon estimates *must* be maintained. Failure to do so causes the filter to believe that it has more information than is really available and, as demonstrated in [2],

this leads to an erroneous map and an inconsistent vehicle estimate.

It has been widely assumed that the stochastic mapping approach of [1] is theoretically sound. Consequently, most recent research has focused on practical issues such as reducing the computational resources, which scale quadratically with map size, required by the optimal algorithm [3–6].

In this paper we re-examine the properties of the stochastic map and consider the problem of a vehicle which possesses a sensor which is capable of measuring range and bearing to beacons in the environment. We show that, when the vehicle is stationary and no process noise acts on it, the joint system is guaranteed to be inconsistent. Furthermore, simulation studies show that a moving vehicle with process noise exhibits a similar type of behavior. However, the time required for the map to become visibly inconsistent is of the order of several hundred timesteps, which is longer than most experimental runs presented in the literature.

The structure of this paper is as follows. The full covariance SLAM algorithm is described in Section 2. The behavior of a stationary vehicle is examined in Section 3 and we derive a theory of a necessary condition which must be met by a map building algorithm. Section 4 describes a simple system and shows that it does not build a consistent map for either a stationary or moving vehicle. Conclusions are drawn in Section 5.

2 Simultaneous Localization and Map Building

The structure of a joint vehicle-beacon system is as follows. The state of the vehicle at timestep k is $\mathbf{x}_v(k)$ and the state of the i th beacon is $\mathbf{p}_i(k)$. The complete state space for a system which comprises of n beacons is

$$\mathbf{x}_n(k) = [\mathbf{x}_v^T(k) \ \mathbf{p}_1^T(k) \ \dots \ \mathbf{p}_n^T(k)]^T. \quad (1)$$

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2001		2. REPORT TYPE		3. DATES COVERED 00-00-2001 to 00-00-2001	
4. TITLE AND SUBTITLE A Counter Example to the Theory of Simultaneous Localization and Map Building			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Virtual Reality Laboratory, Code 5580, ITT Systems/Naval Research Laboratory, Washington, DC, 20375-5320			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES International Conference on Robots and Automation (ICRA) 2001, October 2001, Seoul, South Korea.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

The mean and covariance of this estimate are

$$\hat{\mathbf{x}}_n^T(k|k) = [\hat{\mathbf{x}}_v^T(k|k) \dots \hat{\mathbf{p}}_n^T(k|k)]^T \quad (2)$$

$$\mathbf{P}_n(k|k) = \begin{pmatrix} \mathbf{P}_{vv}(k|k) & \mathbf{P}_{v1}(k|k) & \dots & \mathbf{P}_{vn}(k|k) \\ \mathbf{P}_{1v}(k|k) & \mathbf{P}_{11}(k|k) & \dots & \mathbf{P}_{1n}(k|k) \\ \mathbf{P}_{2v}(k|k) & \mathbf{P}_{21}(k|k) & \dots & \mathbf{P}_{2n}(k|k) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{nv}(k|k) & \mathbf{P}_{n1}(k|k) & \dots & \mathbf{P}_{nn}(k|k) \end{pmatrix} \quad (3)$$

where $\mathbf{P}_{vv}(k|k)$ is the covariance of the AGV's position estimate, $\mathbf{P}_{ii}(k|k)$ is the covariance of the position estimate of the i th beacon and $\mathbf{P}_{ij}(k|k)$ is the cross-correlation between the estimate of i and j .

By assumption, all of the beacons are stationary and no process noise acts upon them. Therefore, the process model is of the form

$$\mathbf{f}[\mathbf{x}_n(k), \mathbf{u}(k), \mathbf{v}(k)] = \begin{pmatrix} \mathbf{f}_v[\mathbf{x}_v(k), \mathbf{u}(k), \mathbf{v}(k)] \\ \mathbf{0} \end{pmatrix} \quad (4)$$

where $\mathbf{u}(k)$ is the control input and $\mathbf{v}(k)$ is the process noise. Similarly, the process noise covariance matrix is

$$\mathbf{Q}_n(k+1) = \begin{bmatrix} \mathbf{Q}_v(k+1) & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix}. \quad (5)$$

The observation model and observation model Jacobian for when the vehicle observes the i th beacon is

$$\mathbf{z}_i(k) = \mathbf{h}_i[\mathbf{x}_v(k), \mathbf{p}_i(k), \mathbf{w}(k)] \quad (6)$$

$$\nabla \mathbf{h}_i = [\nabla \mathbf{h}_i^x \quad -\nabla \mathbf{h}_i^p] \quad (7)$$

where $\mathbf{w}(k)$ is the observation noise with covariance $\mathbf{R}(k)$ and the negative sign on $\nabla \mathbf{h}_i^p$ denotes the fact that the beacon estimates often enter with the opposite sign to the vehicle estimates. Whenever the vehicle observes a beacon which is in the map, the joint system is updated using the Kalman filter update rule¹. If the beacon is not in the map, a new beacon estimate is created and inserted into the map. To initialise the beacon position, the inverse of the beacon observation equation is used:

$$\mathbf{p}_i(k) = \mathbf{g}_i[\mathbf{x}_v(k), \mathbf{z}_i(k), \mathbf{w}(k)].$$

¹These are

$$\hat{\mathbf{x}}(k+1|k+1) = \hat{\mathbf{x}}(k+1|k) + \mathbf{W}(k+1)\nu(k+1), \quad (8)$$

$$\mathbf{P}(k+1|k+1) = \mathbf{P}(k+1|k) - \mathbf{W}(k+1)\mathbf{S}(k+1)\mathbf{W}^T(k+1) \quad (9)$$

$$\nu(k+1) = \mathbf{z}(k+1) - \hat{\mathbf{z}}(k+1|k) \quad (10)$$

$$\mathbf{W}(k+1) = \mathbf{P}(k+1|k)\nabla \mathbf{h}^T\mathbf{S}^{-1}(k+1) \quad (11)$$

$$\mathbf{S}(k+1) = \nabla \mathbf{h}\mathbf{P}(k+1|k)\nabla \mathbf{h}^T + \mathbf{R}(k+1) \quad (12)$$

The new beacon estimate is appended to the map, and the new mean and covariance become:

$$\hat{\mathbf{x}}_{n+1}(k|k) = \begin{bmatrix} \hat{\mathbf{x}}_n(k|k) \\ \mathbf{g}_n[\hat{\mathbf{x}}_v(k|k), \mathbf{z}_n(k), \mathbf{0}] \end{bmatrix}, \quad (13)$$

$$\mathbf{P}_{n+1}(k|k) = \begin{pmatrix} \mathbf{P}_n(k|k) & \mathbf{P}_n(k|k)\nabla \mathbf{g}_n^T \nabla \mathbf{g}_n^x \mathbf{R}(k)\nabla \mathbf{g}_n^w \\ \nabla \mathbf{g}_n^x \mathbf{P}_n(k|k) & \nabla \mathbf{g}_n^x \mathbf{P}_n(k|k)\nabla \mathbf{g}_n^T + \nabla \mathbf{g}_n^w \mathbf{R}(k)\nabla \mathbf{g}_n^w \end{pmatrix}. \quad (14)$$

3 Mapping From A Stationary AGV

Consider the following scenario. An AGV is placed at an unknown location in its environment with a specified mean $\hat{\mathbf{x}}_v(0|0)$ and covariance $\mathbf{P}_{vv}(0|0)$. The AGV then uses a sensor (such as a laser range finder) to measure the position of the beacons relative to the AGV. The AGV is assumed to remain stationary, so no process noise is injected as beacons are initialised and updated. Because the AGV's state is unchanging, and because beacon positions are only measured relative to the AGV, the AGV's position estimate should not change. Therefore, $\hat{\mathbf{x}}_v(k|k) = \hat{\mathbf{x}}_v(0|0)$ for all timesteps k . This condition is satisfied if the following theorem holds:

Theorem 1. *If an AGV estimate is initialised with a non-zero covariance, beacon estimates are initialised using Equations 13 and 14, and all observation covariances are finite, then the state estimate of the AGV will remain unchanged if and only if*

$$\nabla \mathbf{h}_1^x - \nabla \mathbf{h}_1^p \nabla \mathbf{g}_1^x = \mathbf{0} \quad (15)$$

for all timesteps k .

Proof. Assume that the beacon is initialised at timestep 0. We prove Equation 15 by considering the first two timesteps.

The beacon estimate is initialised into the map at timestep 0. Therefore, the state of the system with the initialised beacon is given by Equations 13 and 14. Because the vehicle is stationary and no process noise is injected, the predicted state of the vehicle at timestep 1 is $\hat{\mathbf{x}}_1(1|0) = \hat{\mathbf{x}}_1(0|0)$ and $\mathbf{P}_1(1|0) = \mathbf{P}_1(0|0)$.

From Equation 9, a necessary and sufficient condition to ensure that the vehicle estimate does not change is that the Kalman weight (gain matrix) which is applied to it should be $\mathbf{0}$. Therefore, the weight should be of the form

$$\mathbf{W}(k+1) = \begin{bmatrix} \mathbf{W}_v(k+1) \\ \mathbf{W}_p(k+1) \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{W}_p(k+1) \end{bmatrix},$$

where the dimension of $\mathbf{0}$ is the same as $\hat{\mathbf{x}}_v(k|k)$ and $\mathbf{W}_p(k+1)$ is the weight applied to the beacon. From

Equation 11, the condition is equivalent to demanding that

$$\mathbf{P}(k+1|k) \nabla^T \mathbf{h}_1^x = \begin{bmatrix} \mathbf{0} \\ \mathbf{W}_p \mathbf{S}(k+1) \end{bmatrix}, \quad (16)$$

where the second term is the weight applied to the beacon (not considered here). The actual weight applied to the vehicle is

$$\mathbf{W}_v(1) = \mathbf{P}_v(1|0) (\nabla \mathbf{h}_1^x - \nabla \mathbf{h}_1^p \nabla \mathbf{g}_1^x)^T \mathbf{S}^{-1}(1). \quad (17)$$

Because $\mathbf{P}_v(1|0)$ and $\mathbf{S}^{-1}(1)$ are nonsingular, a necessary and sufficient condition for $\mathbf{W}_v(1)$ to be $\mathbf{0}$ is

$$\mathbf{0} = \nabla \mathbf{h}_1^x - \nabla \mathbf{h}_1^p \nabla \mathbf{g}_1^x.$$

This result proves the theorem for a single timestep. In [7] we show that it can be readily extended to an arbitrary number of timesteps. \square

The importance of this result is that the behavior of the vehicle and beacon depends critically on the Jacobian matrices used to initialise the beacon. In the special case that the observation model is linear and time invariant² ($\nabla \mathbf{h}_i^x = \mathbf{H}^x$, $\nabla \mathbf{h}_i^p = \mathbf{H}^p$ and $\nabla \mathbf{g}_i^x = \mathbf{G}^x$), the condition requires that $\mathbf{H}^p \mathbf{G}^x = \mathbf{H}^x$. However, in a general nonlinear system where the Jacobian matrices are functions of noisy observations and erroneous estimates, it is not clear that the condition in Equation 15 can be guaranteed to hold. Furthermore, because this condition is a *structural* relationship, normal tuning procedures (such as inflating the observation noise covariances) cannot circumvent the problem. If the vehicle estimate changes for one value of $\mathbf{R}(k)$, it will change for all (finite) values of $\mathbf{R}(k)$. In fact, as we now show, this condition cannot be guaranteed even for the simplest case of a position-orientation AGV model and a range-bearing sensor model.

4 A Concrete Example

Consider the following simple system. The vehicle state is described by its position (x_v, y_v) and orientation θ_v in some global coordinate system. The vehicle is equipped with a sensor that is able to return the range r and bearing ϕ of a target relative to the sensor platform. The sensor is a rotating scanner (such as a laser range finder) which completes one revolution per second.

²Even with nonlinear process and observation models, such a system structure arises in the relative map [6] and geometric projection filter [5] which decouple the problem of map building from vehicle localization. However, with these approaches it appears that the map can only be used as an aid to beacon gating and it cannot be used to update the vehicle position estimate.

4.1 System Equations

The vehicle model is the standard equation for a steered bicycle [8]:

$$\mathbf{x}_v(k+1) = \begin{pmatrix} x_v(k) + V(k+1)\Delta T \cos(\delta(k+1) + \theta_v(k)) \\ y_v(k) + V(k+1)\Delta T \sin(\delta(k+1) + \theta_v(k)) \\ \theta_v(k) + \frac{V(k+1)\Delta T \sin(\delta(k+1))}{B} \end{pmatrix},$$

where the timestep is ΔT , the control inputs are the wheel speed $V(k+1)$ and steer angle $\delta(k+1)$ and the vehicle wheel base is B . The process noises are additive disturbances which act on $V(k)$ and $\delta(k)$. The observation model is

$$\mathbf{h}_i[\mathbf{x}_v(k), \mathbf{p}_i(k), \mathbf{w}(k)] = \begin{bmatrix} \sqrt{(x_i - x_v)^2 + (y_i - y_v)^2} \\ \tan^{-1}\left(\frac{y_i - y_v}{x_i - x_v}\right) - \theta_v \end{bmatrix} \quad (18)$$

The Jacobian for this equation can be written as

$$\nabla \mathbf{h}_i^x = (\mathbf{h}_i^{x_v y_v} \quad \mathbf{h}_i^\phi \quad \mathbf{h}_i^{x_i y_i}),$$

where

$$\mathbf{h}_i^{x_v y_v} = \begin{pmatrix} -(x_i - x_v)/r & -(y_i - y_v)/r \\ (y_i - y_v)/r^2 & -(x_i - x_v)/r^2 \end{pmatrix}, \quad (19)$$

$$\mathbf{h}_i^\phi = \begin{bmatrix} 0 \\ -1 \end{bmatrix}, \quad (20)$$

$$\mathbf{h}_i^{x_i y_i} = -\mathbf{h}_i^{x_v y_v}. \quad (21)$$

The observation noise Jacobian is

$$\mathbf{h}_i^w = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (22)$$

Let $\alpha(k) = \theta_v(k) + \phi(k)$. Inverting Equation 18, the beacon position is initialised as

$$\mathbf{g}_i[\mathbf{x}_v(k), \mathbf{w}(k)] = \begin{bmatrix} x_v(k) + r(k) \cos[\alpha(k)] \\ y_v(k) + r(k) \sin[\alpha(k)] \end{bmatrix}. \quad (23)$$

and the Jacobian is

$$\nabla \mathbf{g}_i^x = (\mathbf{g}_i^{x_v y_v} \quad \mathbf{g}_i^\theta),$$

where

$$\mathbf{g}_i^{x_v y_v} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (24)$$

$$\mathbf{g}_i^\theta = \begin{bmatrix} -r(k) \sin[\alpha(k)] \\ r(k) \cos[\theta_v(k) + \phi(k)] \end{bmatrix}, \quad (25)$$

$$\mathbf{g}_i^w = \begin{bmatrix} \cos[\alpha(k)] & -r(k) \sin[\alpha(k)] \\ \sin[\alpha(k)] & r(k) \cos[\alpha(k)] \end{bmatrix}. \quad (26)$$

We now consider two special cases for this system — a stationary vehicle with no process noise, and a vehicle which moves in a circle with nominally constant control inputs.

4.2 Behavior of a Stationary Vehicle

First consider the special case that the vehicle is stationary, there is only one beacon, and *no process noise* acts on the system. In this situation the process model reduces to the identity matrix and

$$\mathbf{f}_v[\mathbf{x}_v(k), \mathbf{0}, \mathbf{0}] = \mathbf{x}_v(k).$$

where the dimension of $\mathbf{0}$ is the same as $\hat{\mathbf{x}}_v(k|k)$. At timestep 0 the beacon is initialised into the map and the mean and covariance are set using Equations 13 and 14. This case can be analysed by Theorem 1. Assuming that Equation 16 has been obeyed up to timestep k , the numerator of $\mathbf{W}(k)$ is

$$\begin{pmatrix} \mathbf{P}_n(k|k) & \mathbf{P}_n(k|k) \nabla \mathbf{g}_n^T \mathbf{g}_n^x \\ \nabla \mathbf{g}_n^x \mathbf{P}_n(k|k) & \nabla \mathbf{g}_n^x \mathbf{P}_n(k|k) \nabla \mathbf{g}_n^T \mathbf{g}_n^x + \mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{h}_i^{x_v y_v T} \\ \mathbf{h}_i^{\phi T} \\ \mathbf{h}_i^{x_i y_i T} \end{pmatrix}$$

where $\mathbf{A}(k)$ is a positive semidefinite matrix corresponding to the partially filtered observation noise. Therefore, the weight on the vehicle \mathbf{W}_v is

$$\mathbf{W}_v = \mathbf{P}_n(k|k) \left\{ \begin{bmatrix} \mathbf{h}_i^{x_v y_v T} \\ \mathbf{h}_i^{\theta T} \end{bmatrix} + \begin{bmatrix} \mathbf{g}_i^{x_v y_v T} \\ \mathbf{g}_i^{\theta T} \end{bmatrix} \mathbf{h}_i^{x_i y_i T} \right\}.$$

Substituting from Equations 24 to 26 and using the property that $\mathbf{h}_i^{x_v y_v T} = -\mathbf{h}_i^{x_i y_i T}$,

$$\mathbf{W}_v = \mathbf{P}_n(k|k) \begin{bmatrix} \mathbf{0} \\ \mathbf{h}_i^{\theta T} - \mathbf{g}_i^{\theta T} \mathbf{h}_i^{x_v y_v T} \end{bmatrix}.$$

Therefore, the weight on the vehicle position states is always guaranteed to be $\mathbf{0}$. However, this is *not* the case for the vehicle orientation state. Substituting from Equations 20, 25 and 19, it can be readily shown that this weight is non-zero only if the angle between the vehicle and the beacon ($\theta_v + \phi$) is the *same* as the value when the beacon was first initialised [7]. It should be emphasised that these analytical results reflect a fundamental failure in the structure of the cross correlation between the vehicle and beacon estimates. The errors occur irrespective of the magnitude of covariances and they are not the result of subtle numerical implementation errors.

This behavior can be clearly seen in a simulation study of this scenario. For the results in this paper we use the following conditions. The vehicle initially starts at the origin with the standard deviations in x_v and y_v of 0.7m and θ_v of 5° . It observes a beacon at (97.89, 70.1) with an accurate sensor whose observation noise covariance is $\mathbf{R}(k) = \text{diag}(0.25\text{m}^2, (1^\circ)^2)$. Figure 1 plots the time history of the estimates of x_v , θ_v , x_1 and y_1 . As expected, the estimate of x_v does

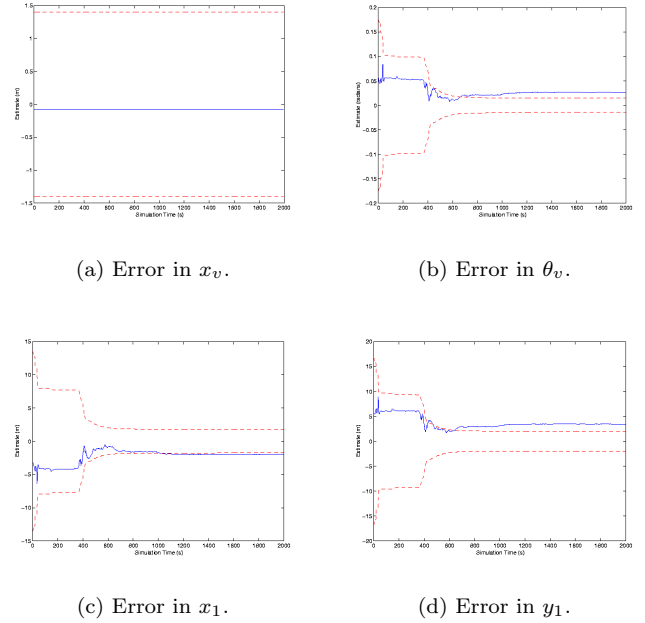


Figure 1: Estimation errors and 2 standard deviation bounds. The standard deviation bounds are shown as dashed lines.

not change. However, the estimate of θ_v immediately starts to change and its covariance begins to decline. By the end of the run, the orientation covariance is less than 6% of its initial value. However, this update is entirely spurious — there is no additional information about the vehicle’s orientation and so the orientation estimate becomes inconsistent. In turn, the beacon estimate becomes inconsistent as well. Extending the simulation further shows that, in the limit, the steady-state covariance of the beacon estimate does not become 0 but is a value greater than the initial vehicle position covariance.

The value of $(\theta_v + \phi)$ also changes if the vehicle and beacon configuration changes. Figure 2 shows the result when the vehicle, at timestep 10, “teleports” to the position (50, 50). This jump occurs instantaneously and without uncertainty, i.e., at the same time the estimate changes, the true vehicle position changes by the same amount³. As can be seen the results are dramatic. Although the error in the estimate of x_v remains unchanged after the instantaneous translation, observations of the beacons lead to a large spurious reduction in both the vehicle orientation and the beacon position estimates. Again, these reductions correspond to inconsistent estimates.

³This is equivalent to the process model $x_v(k+1) = x_v(k) + \Delta x$, $y_v(k+1) = y_v(k) + \Delta y$ where Δx and Δy are known.

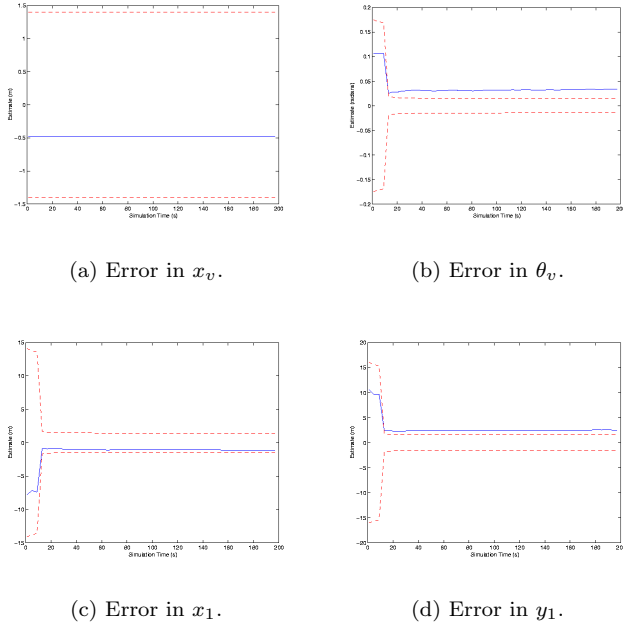


Figure 2: Estimation errors and 2 standard deviation bounds. For a vehicle which “teleports” from position (0,0) to (50,50) at timestep 10.

It must be emphasised that the inconsistencies in both examples are not simply due to the fact that the observation Jacobian is calculated using the noise-corrupted sensor observations. Even if the *true* state of the beacon and the vehicle were always used to calculate the Jacobian of the observation equation, *any* motion by the vehicle that affects the observation Jacobian will lead to inconsistency analogous to a violation of the condition of Equation 15 in the stationary case described in Section 3. In fact, a perfectly known change in the AGV’s orientation – *with no change in position* – will have the same effect. In the next section we generalise the example to include the accumulation of process noise by a moving vehicle.

4.3 The Behavior of a Moving Vehicle

In this section we consider the case of a moving vehicle in a long-duration SLAM simulation. As in our example of a stationary vehicle with a range-bearing sensor, our goal is to keep the scenario as simple as possible to demonstrate the perniciousness of the inconsistency problem. To this end, we assume that the vehicle travels in a circle with constant control inputs $\mathbf{u}(k+1) = [1 \ 2^\circ]$ and process noise standard deviations $\mathbf{Q}(k) = \text{diag}\{0.25, (0.3^\circ)^2\}$. The vehicle observes an environment which contains 5 beacons.

The time history of the first 600 timesteps of the vehicle and beacon estimates are shown in Figure 3 and Figure 4 respectively. The errors in the beacons and the beacon estimates appear to have stabilised within the two sigma bounds and it appears that this algorithm is performing in a satisfactory manner.




Figure 3: Results for the first 600 seconds (3000 timesteps) of a moving vehicle.




Figure 4: Results for the first 600 seconds (3000 timesteps) of beacon 1.

Figures 5 and 6 show the time histories when the experiment is allowed to run for 16,000 timesteps. These results clearly show that the apparent consistency is only a short-term phenomena: within less than five thousand time steps (1000 updates per beacon) the map has become inconsistent. We speculate that this behavior has not been recognised in the literature for two main reasons. First, most systems reported in the literature (such as [9] or [6]) only present results for the first few hundred timesteps. Over such short durations the signs of divergence are not very prominent. Second, the few long-duration studies (such as [3]) have all used compasses to measure the absolute orientation of the vehicle. A compass causes the orientation errors to be filtered out, and significantly reduces the rate of divergence [7].




Figure 5: Results for the first 3200 seconds (16000 timesteps).




Figure 6: Estimation errors and 2 standard deviation bounds for beacon 1 for 3200 seconds (16000 timesteps). The standard deviation bounds are the pairs of dashed lines.

5 Conclusions

This paper has analysed the properties of the full-covariance stochastic mapping approach to SLAM. We have shown that a simple but realistic scenario is guaranteed to produce inconsistent vehicle and beacon estimates. Furthermore, we have shown that the problem typically occurs after several hundred timesteps, which is longer than the duration of many experimental runs. The full implications of this result (for example, whether it is possible to build a consistent map without adding process noise to the beacon estimates) is presently under investigation [7]. However, two conclusions can be immediately drawn from these results. First, it is questionable whether the Kalman filter framework developed in [1] provides a general, robust and rigorous solution to the stochastic SLAM problem. Second, the consistency of *any* map building algorithm cannot be fully assessed from experimental runs which are of short duration. As demonstrated

in Subsection 4.3, an algorithm can appear to be consistent for many hundreds of time steps but, in fact, might prove to be inconsistent.

References

- [1] R. Smith, M. Self and P. Cheeseman, *Estimating Uncertain Spatial Relationships in Robotics, Autonomous Robot Vehicles*. Springer-Verlag, 1989.
- [2] J. A. Castellanos, J.D. Tardós and G. Schmidt, “Building a Global Map of the Environment of a Mobile Robot: The Importance of Correlations,” in *Proceedings of 1997 IEEE Conference on Robotics and Automation*, pp. 1053–1059, IEEE Press, April 1997.
- [3] J. J. Leonard and H. J. S. Feder, “A computationally efficient method for large-scale concurrent mapping and localization,” in *Ninth International Symposium on Robotics Research* (D. Koditschek and J. Hollerbach, eds.), (Snowbird, Utah), October 1999.
- [4] J. K. Uhlmann, S. J. Julier and M. Csorba, “Non-divergent Simultaneous Map Building and Localisation using Covariance Intersection,” in *The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Orlando, Florida, SPIE, 1997*. Navigation and Control Technologies for Unmanned Systems II.
- [5] P. Newman, *On the Structure and Solution of the Simultaneous Localisation and Map Building Problem*. PhD thesis, University of Sydney, 1999.
- [6] M. Csorba and H. F. Durrant-Whyte, “New Approach to Map Building Using Relative Position Estimates,” in *Proceedings of the SPIE AeroSense Conference, Orlando, FL*, vol. 3087, pp. 115–125, April 1994.
- [7] S. J. Julier and J. K. Uhlmann, “Fundamental Theoretical Limitations in the Theory of Simultaneous Localisation and Map Building.” In preparation, 2000.
- [8] H. F. Durrant-Whyte, “An autonomous guided vehicle for cargo handling applications,” *International Journal of Robotics Research*, vol. 15, pp. 407–440, October 1996.
- [9] J. A. Castellanos, J. M. M. Montiel, J. Neira and J. D. Tardós, “The SPmap: A Probabilistic Framework for Simultaneous Localization and Map Building,” *IEEE Transactions on Robotics and Automation*, vol. 15, pp. 948–952, October 1999.